

Cloud-Assisted Remote Sensor Network Virtualization for Distributed Consensus Estimation

Sherif Abdelwahab*, Bechir Hamdaoui*, and Mohsen Guizani†

* Oregon State University, abdelwas,hamdaoui@eecs.oregonstate.edu

† Qatar University, mguizani@ieee.org

Abstract

We develop cloud-assisted remote sensing techniques for enabling distributed consensus estimation of unknown parameters in a given geographic area. We first propose a distributed sensor network virtualization algorithm that searches for, selects, and coordinates Internet-accessible sensors to perform a sensing task in a specific region. The algorithm converges in linearithmic time for large-scale networks, and requires exchanging a number of messages that is at most linear in the number of sensors. Second, we design an uncoordinated, distributed algorithm that relies on the selected sensors to estimate a set of parameters without requiring synchronization among the sensors. Our simulation results show that the proposed algorithm, when compared to conventional ADMM (Alternating Direction Method of Multipliers), reduces communication overhead significantly without compromising the estimation error. In addition, the convergence time, though increases slightly, is still linear as in the case of conventional ADMM.

1 Introduction

As the Internet of Things (IoT) emerges, a rapid growth of the number of sensor-equipped things (e.g., smart-phones, tablets, etc.) is observed where it became important to rethink the way conventional sensor-based management techniques are designed. Although sensor-based distributed sensing has already been investigated in the past, cloud-assisted remote sensing is a new paradigm that capitalizes on the capabilities of IoT to enable what is called *Sensing as a Service* [1]. Cloud-assisted remote sensing based distributed parameter estimation is one example of such services.

Traditional sensor networks depend primarily on the sophistication and accuracy of the sensory devices themselves to perform sensing tasks and meet quality of service requirements. In the Sensing as a Service model, cloud-based sensor networks rely mainly on swarms of participatory sensors to perform remote sensing tasks. Unlike traditional sensory devices, participatory sensors, though come with new opportunities, present key challenges, mainly pertaining to their sporadic availability and unpredictable mobility.

In the cloud-based remote sensing paradigm, a cloud agent (or manager) is responsible for receiving and handling

remote sensing task requests from cloud clients. When a request is granted, the agent is also responsible for virtualizing a sensor network to perform the requested sensing task. An example of such tasks is distributed consensus estimation, in which an unknown set of parameters need to be estimated. Imagine for example the virtual sensing task in which it is required to track the location of an RFID-tagged person in a large campus. A group of smart-phones connected through machine-to-machine physical links can estimate the location of a person based on the RFID signal strength each smart-phone receives and measures. In this case, a virtual sensing request is sent to the cloud agent which dispatches the request to few smart-phones in the swarm. These smart-phones autonomously search for a group of smart-phones that can read RFID tags and are willing to participate in the sensing task to estimate the distance based on received signal strengths. Among all these smart-phones, a subgroup of them is then selected to form a virtual sensor network connected according to a topology that is to be specified by the cloud agent itself. Such virtual sensor network distributedly and cooperatively estimates the location of the person using distributed linear consensus algorithms (see for example [20, 11]), and continuously sends an update of the location of the person to the cloud agent, which in turn sends it to the cloud client.

In this paper, we develop cloud-assisted remote sensing algorithms that enable distributed consensus estimation of unknown parameters. Specifically, we propose:

- An efficient network virtualization algorithm that can search and select sensors from the swarm to form a virtual sensor network that can perform a sensing task. The algorithm consists of selecting a set of sensors that are willing to participate in the requested remote sensing task, and finding optimal one-to-one mappings between virtual and participatory sensors.
- An efficient estimation algorithm that relies on the virtual sensor network, formed by our proposed virtualization algorithm, to estimate a set of unknown parameters in a distributed way and without requiring synchronization among participatory sensors.

Our results show that given a virtual sensing task requiring g sensors and a swarm of n sensors where a pair of sensors is considered connected if the sensors are within a

distance r from one another, our virtualization algorithm finds the capabilities of the sensors in $O(r^{-1} \log n)$ with an average number of messages per sensor of $\Theta(1)$. Our results also show that our proposed algorithms achieve a virtualization benefit that is very close to an upper bound in $O(\max\{r^{-1}n \log n, g^3\})$ time with a $O(n)$ worst case average number of messages per sensor. Finally, our results show that our proposed distributed parameter estimation algorithm has a linear convergence time and incurs communication overhead that is at least an order of magnitude lesser than that incurred by the conventional Alternating Direction Method of Multipliers (ADMM).

In our design, the cloud agent does not need to have full knowledge of the underlying swarm of sensors, including its topology, sensor capabilities, and sensor availability. We only assume that the cloud agent has a direct communication with some, but not necessarily all, sensors. Each sensor has to exchange small sized messages with only its direct neighbors to diffuse information across the whole or part of the swarm. We also assume that the swarm is dynamic by nature in that sensors' availability, locations, connectivity, and capacities may change sporadically and unpredictably over time.

This paper is organized as follows: In Section 2, we overview our algorithmic design, describe the requirements, and state the related work. In Section 3, we propose and present our sensor network virtualization algorithm. In Section 4, we propose and present our distributed parameter estimation algorithm. Finally, we evaluate and compare the performance of our proposed algorithms numerically in Section 5 and conclude the paper in Section 6.

2 Framework Overview: System Model, Problem Formulation, and Related Work

In this section, we provide a brief overview of the different framework components and algorithms that are proposed in this paper to perform cloud-assisted remote sensing based parameter estimation. We also talk about some previously done works that are related to the algorithms proposed in this paper.

2.1 Participatory Sensors

We consider a swarm of a large number of participatory sensors that are managed by a cloud platform. Each sensor of the swarm is cloud accessible via the Internet and is assumed to have some sensing capability. We model the swarm as an Euclidean geometric random graph $\mathcal{G} = (S, L)$, where S is a set of n sensors, and L is the set of all links connecting the sensors, where two sensors are considered to be connected if they are within a transmission radius, r , of each other. Let $loc(i)$ denote the physical location of sensor i , and $C(i)$ denote its sensing capability or capacity ($C(i)$ can for e.g. refer to maximum allowed sensing time, maxi-

mum allowed processing power, maximum allowed memory capacity, etc.).

We assume that each sensor $i \in S$ is capable of estimating a vector of unknown parameters, $\theta \in \mathbf{R}^N$, through noisy measurements, $x_i \in \mathbf{R}^M$. That is,

$$x_i = H_i \theta + u_i, \quad i = 1, \dots, n$$

where $H_i \in \mathbf{R}^{M \times N}$ is sensor i 's sensing model (typically known to i only) relating x_i to θ , and u_i is an additive Gaussian noise with zero mean and variance σ_i^2 . We assume that u_i and u_j are independent from one another for all $i, j \in S$. Because different sensors may have different sensing models and/or different measurement methods, it is very likely that different sensors have different estimates of θ . Also, we do not assume/require that the sensors are synchronized; that is, the consensus algorithms we develop in this paper to estimate θ are asynchronous.

2.2 Virtual Remote Sensing

We assume that there exists a cloud manager/agent that is responsible for managing the participatory sensors, handling virtual sensing task requests (to be submitted by cloud clients), and ensuring that clients' Service Level Agreements (SLAs) are met once their requests are granted by the cloud. Each virtual sensing task request is represented by a quadruple, $\langle g, \theta, c, \delta \rangle$, where g is the number of (virtual) sensors requested to perform the sensing task, $\theta \in \mathbf{R}^N$ is a column vector of unknown parameters to estimate by the virtual sensors, and the location c and radius δ indicate the area of interest that needs to be sensed; i.e., all requested sensors must be located within δ distance from the center c . Generally speaking, SLAs consist of: (i) a maximum time within which the sensing task must be completed, (ii) an absolute tolerance $\epsilon_{\text{abs}} > 0$ of the estimation quality, (iii) a relative tolerance $\epsilon_{\text{rel}} > 0$ of the estimation quality (maximum gap between the g sensors' local estimates of θ), and (iv) a maximum rejection rate, defined as the ratio of the number of failed virtualizations to the total number of sensing task requests.

Upon receiving a sensing task request, the cloud agent's job is then to define a set V of g virtual sensors to be realized by g connected participatory sensors, all located within distance δ from the center c , that can collaboratively and distributively estimate θ . Depending on the SLAs of the sensing task, the cloud agent needs to determine each of the following three parameters before proceeding with sensor network virtualization needed to perform a requested sensing task. First, it needs to choose a suitable virtual topology that connects the set of virtual sensors, V , so that they can perform the sensing task collaboratively. Although other topologies can be used, we focus in this paper on three types: complete, cyclic, and star. For a given topology, let E denote the set of virtual links connecting the virtual sensors and $\Upsilon = (V, E)$ be the graph representing the virtual sensor network. Note that the cloud agent needs to make sure that the virtual sensor network is connected according to the chosen topology.

Another parameter the cloud agent needs to fix and associate with Υ is the maximum allowed path length, \bar{h} , between any pair of virtual sensors. \bar{h} can be viewed as a way to limit the number of sensors/hops message exchanges among virtual sensors can go through. It is a parameter that can be used to impose an upper bound on end-to-end message delays, and thus, on the sensing task completion time. Note that a virtual link between two virtual sensors may be realized by more than two physical sensors, and some of these sensors may not necessarily realize a virtual sensor in V by itself. That is, some sensors may only be used for forwarding traffic without participating in performing the sensing task.

The third parameter the cloud agent needs to choose and set is the sensing capacity threshold $R(j)$ for a virtual sensor $j \in V$. The capacity $C(i)$ of a participatory sensor i which realizes j must then exceeds the capacity threshold $R(j)$. This threshold can for e.g. represents the minimum storage capacity, the minimum CPU computing power, and/or the minimum amount of time required by the sensing task.

The required values of the parameters associated with Υ including E , $R(j)$, and \bar{h} will be mainly determined by the client's SLA, and is beyond the scope of this paper. In this paper, we focus instead on the design of efficient algorithms that meet these design requirements. Our ultimate objective in this work is to design a distributed consensus algorithm that enables the estimation of the unknown parameter vector, θ , subject to the design parameter requirements specified by both the cloud client and the cloud agent. To this end, the key tasks that need to be executed by the swarm of participatory sensors to perform estimation are:

2.2.1 Sensor Search

It consists of searching for the sensors, among all participatory sensors, that meet the Υ requirements. More specifically, the swarm searches for a subset of participatory sensors, $S' \subset S$, such that a sensor $i \in S'$ if:

- i) it can sense and estimate θ through a sensing model H_i ; i.e., it can observe a vector x_i that can be expressed as $x_i = H_i\theta + u_i$,
- ii) it is geographically located within δ distance from c , and
- iii) its capacity $C(i) \geq R(j)$ for at least one virtual sensor $j \in V$.

For each participatory sensor i , we define the virtual domain of i , $\mathcal{D}(i)$, as the set of all virtual sensors that can be supported by sensor i ; that is,

$$\mathcal{D}(i) = \begin{cases} \{j \in V : C(i) \geq R(j)\} & \text{if } \|loc(i) - c\| \leq \delta \\ \emptyset & \text{otherwise,} \end{cases} \quad (1)$$

and the objective of this task is to construct, with the minimum possible communication overhead, each participatory sensor i 's virtual domain, $\mathcal{D}(i)$, and to determine the set S' as fast as possible, all without assuming prior knowledge

of the \mathcal{G} topology. Our proposed technique for performing such a task is presented in Section 3.

Related work. In a recent work, Perera et. al [18, 17] described a system of context-aware sensor search to address the research challenges of searching for sensors when large numbers of sensors with overlapping and redundant functionality are available to the cloud. Such a sensor search approach suffers from practical limitations as it relies on centralized knowledge of all available sensors and requires continuous tracking of the sensors' dynamics, such as the sensors' availability, connectivity, and mobility.

Sensor search algorithms need to be simple, have bounded search latency, and incur minimum communication overhead between the cloud platform and the large number of participatory sensors. Unlike centralized resource discovery algorithms [3], gossip-based search protocols are distributed and topology-independent, which are more suitable for sensor search in the IoT context. Gossip protocols are originally designed for information dissemination [4], and have been demonstrated to be effective in resource discovery in Peer-to-Peer (P2P) networks [9]. Our proposed framework relies on gossip techniques to determine the set of sensors among all participatory sensors that are capable of performing sensing and are willing to participate in the formation of a virtual sensor network.

2.2.2 Sensor Network Virtualization

This virtualization task consists of finding (i) a set $A \subset S'$ of exactly g connected sensors selected among all sensors in S' (the g selected sensors should be connected according to the virtual topology chosen by the cloud agent) and (ii) a set $\mathcal{M}_A \subset \{(i, j) \in A \times V : j \in \mathcal{D}(i)\}$ of one-to-one mapped pairs (each participatory sensor in A is mapped to one and only one virtual sensor in V) such that the length, $h(i, i')$, of any simple path connecting two distinct participatory sensors i, i' in A mapping a pair of directly connected virtual sensors $(j, j') \in E$ is less than or equal to \bar{h} . We refer to a possible $\{A, \mathcal{M}_A\}$ pair as a *feasible virtualization* of the requested virtual sensor network Υ . Note that for any possible set A , there could exit multiple possible sets, \mathcal{M}_A , each can form a feasible virtualization when paired with A , and the objective of a sensor network virtualization algorithm is then to find the 'optimal' feasible virtualization, $\{A, \mathcal{M}_A\}^*$.

We now define and introduce what an 'optimal' feasible virtualization means. We consider that the cost (to the cloud) of virtualizing a virtual sensor network request, $\Upsilon = (V, E)$, is determined by the amounts of requested resources and given by $\text{Cost}(\Upsilon) = \alpha|V| + \beta|E|$, where α denotes an incentive paid by the cloud to each participatory sensor, and β denotes an incentive associated with each physical path between each pair of participatory sensors in A . An incentive could be monetary or could be in any other form (e.g., credits, services, etc.). On the other hand, the total benefit (to the cloud/swarm of sensors) resulting from

a feasible virtualization, $\{A, \mathcal{M}_A\}$, can be expressed as

$$\text{Benefit} = \sum_{(i,j) \in \mathcal{M}_A} \alpha \frac{C(i) - R(j)}{C(i)} + \sum_{(i,i') \in P} \beta \frac{\bar{h} - h(i,i')}{\bar{h}}, \quad (2)$$

where $h(i, i')$ is again the path length (in number of hops) of the path connecting the sensors of the pair (i, i') mapping the virtual link between j and j' and $P = \{(i, i') \in A \times A : (i, j), (i', j') \in \mathcal{M}_A, (j, j') \in E\}$ denotes the set of all such pairs. Note that for a given sensing task request (i.e., for a given virtualization cost), the lesser the used physical resources, the higher the total benefit. The sensor network virtualization algorithm that we propose consists then of finding a feasible virtualization that maximizes the total benefit given in Eq. (2). We refer to the optimal solution as $\{A, \mathcal{M}_A\}^*$. Clearly, finding $\{A, \mathcal{M}_A\}^*$ is a hard problem due to the factorial size of the solution space (in n). Our first contribution in this work is to develop an algorithm that solves this virtualization problem efficiently. The algorithm is presented in Section 3.

Related work. Network virtualization techniques proposed in the past decade consist mainly of virtual network embedding algorithms, which instantiate virtual networks on substrate infrastructures [10, 3]. Most of these virtual network embedding algorithms are centralized (e.g. [7]) due to the ease of deployment of centralized approaches in cloud platforms where the cloud provider desires to have full control on the physical network resources.

Distributed virtual network embedding and virtualization algorithms have also been proposed in literature [12, 2]. One of the limitation of the algorithm proposed in [12] lies in its unsuitability for swarm virtualization, as the authors assume unlimited physical resources and consider an offline resource virtualization approach. As for the more recent work proposed in [2], although the virtualization phase of the physical network does not require full/global knowledge of the swarm, the cloud must initially partition the swarm into hierarchies and delegate each virtual network request to a different hierarchy, which also requires full knowledge but about the sensors in each hierarchy. Unlike this approach, our proposed virtualization algorithm is fully distributed. In addition, it does not require any synchronization among sensors.

2.2.3 Distributed Consensus Estimation

This task relies on the virtual sensor network (formed in the previous step) to provide an estimate of θ that is at most ϵ_{abs} from the optimal, and that all the virtualized sensors consent to the same estimate value of θ with a tolerance of ϵ_{rel} .

Without loss of generality, consider indexing the selected g sensors in the virtual sensor network as $1 \dots g$ and let $x = [x_1^T, \dots, x_g^T]^T$, $H = [H_1^T, \dots, H_g^T]^T$, and $u = [u_1^T, \dots, u_g^T]^T$. The aggregate measurements can then be written as $x = H\theta + u$. One simple approach of estimating θ is to have the cloud agent first collect from each virtual sensor i its

measurement vector, x_i , and its sensing model, H_i , and then solves the following Least Squares (LS) problem

$$\text{minimize} \quad \frac{1}{2} \|x - H\hat{\theta}\|^2 \quad (3)$$

where $\hat{\theta}$ is here the optimization variable. The unbiased maximum-likelihood (ML) estimate of θ is simply $\hat{\theta}_{\text{LS}} = (H^T H)^{-1} H^T x$.

This centralized LS approach, though simple, requires that each virtualized sensor exchanges its measurement vector and its sensing model with the cloud agent, which can create significant communication overhead, especially when the number of measurements, M , and the number of virtual sensors, g , are large.

We instead propose, in this paper, a decentralized approach that relies on the virtual sensor network to provide an estimation of the parameter vector θ . We rely on the recent results presented in [21] to develop our distributed estimation algorithm, which reduces communication overhead significantly when compared to the conventional ADMM approach [16] in addition to not requiring synchronization among sensors. The proposed algorithm is presented in Section 4.

Related work. Distributed parameter estimation approaches have been proposed in [16, 20, 11]. Estimation can for e.g. be carried out by first computing a local estimate at each virtual sensor and then perform a distributed weighted average of the local estimates [20]. This approach results in an ML estimate, but does not limit/bound the variation between mean square errors of local estimates. More recently, Paul et al. [16] propose a distributed estimation algorithm based on ADMM. Although this approach results in an optimal mean square error when compared to LS, it exhibits a significant in-network communication overhead that requires even more messages to be exchanged among sensors than that exchanged in the centralized LS. One approach also proposed in [16] to overcome this problem is to approximate the computation of primal and dual variables at each step of the algorithm by using earlier versions of these variables instead of sharing them at each iteration which marginally reduces the communication overhead. In addition to the increased communication overhead, conventional ADMM requires synchronous operation of the sensors. This is very challenging from a practical viewpoint, and does not scale well especially when applied in the IoT context. It has been shown recently that an asynchronous implementation of ADMM has $O(1/k)$ convergence [21].

Not only is our proposed estimation algorithm both asynchronous and distributed, but also reduces communication overhead significantly when compared to the conventional ADMM approach [16].

3 Sensor Network Virtualization

We begin by presenting our proposed Randomized and Asynchronous Distributed Virtualization (RADV) algorithm, which consists of four phases: (I) searching for sensors that can support a virtual sensing task request Υ , (II)

pruning of virtual domains $\mathcal{D}(i)$ for all $i \in S$, (III) construction of benefit matrices in a distributed manner, and (IV) solving assignment problems at virtual sensors. This approach results in multiple solutions each evaluated by a different sensor, and the cloud agent selects the solution with the maximum benefit.

We design time-invariant gossip based algorithms for the first three phases in which sensors exchange information randomly [19, 4]. At the k -th time slot, let sensor i be active and contact its neighbor sensor j (i.e., $(i, j) \in L$) with probability $T_{i,j} > 0$ only if j can be contacted by more than one of its neighbors. The probability $T_{i,i}$ denotes the probability that i does not contact any other sensor. Let the $n \times n$ matrix $T = [T_{i,j}]$ be a doubly stochastic transition matrix of non-negative entries [5]. A natural choice of $T_{i,j}$ is

$$T_{i,j} = \begin{cases} \frac{1}{d_i + 1}, & \text{if } i = j \text{ or } (i, j) \in L, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where $d_i = |\{j \in S : (i, j) \in L\}|$ is the degree of sensor i .

We now present each of the four phases in greater details.

Phase I—Sensor Search

The objective of this phase is to construct $\mathcal{D}(s)$ for all $s \in S$ and S' as fast as possible without assuming prior knowledge of the \mathcal{G} topology. First, the cloud initiates the search at time $k = 0$ by sending Υ to one or more arbitrary sensors. At any later time k , s and its neighbor s' exchange information as follows. s pushes Υ to s' only if s' does not have Υ , or pulls it from s' only if s does not have Υ . If s contacts s' and both s and s' have received Υ before, s stops contacting any other sensor. Upon receipt of Υ , a sensor s constructs $\mathcal{D}(s)$ according to Eq.(1) and S' as

$$S' = \{s \in S : |\mathcal{D}(s)| > 0\}.$$

A virtual domain of sensor i , $\mathcal{D}(i)$, evaluated during the sensor search phase is not sufficient to tell whether a feasible embedding can be found if sensor i virtualizes (is mapped to) j . For example, the virtualization in which sensor i virtualizes j and sensor i' virtualizes j' where there is a virtual link (j, j') is not feasible if i' is not reachable from i and vice versa. When this situation occurs, we say that sensor i is incapable of supporting the topology requirement specified by E , initiating thus a virtual domain pruning.

Phase II—Virtual Domain Pruning

During this phase, we ensure that all virtualized sensors maintain the topology E by allowing a sensor to receive the virtual domains of other sensors and delete a virtual sensor j from its domain if there exists a virtual link (j, j') such that j' is not included in any other received domains. Let $D_s \subset \{\mathcal{D}(i) : i \in S\}$ denotes the set of domains that sensor s has at time k . Initially $D_s = \{\mathcal{D}(s)\}$ and $h(i, s) = 0$ for all $i \in S^1$. Using the same transition matrix, T , defined

¹Knowledge about other sensors existence is not needed, and h is typically evaluated dynamically.

in Eq. (4), s contacts only one of its neighbors s' at time k . Then, for all $\mathcal{D}(i) \in D_s : i \neq s'$, s pushes $\mathcal{D}(i)$ to s' only if s' did not receive $\mathcal{D}(i)$ before and $h(i, s) < \bar{h}$. Also, for all $\mathcal{D}(i) \in D_{s'} : i \neq s$, s pulls $\mathcal{D}(i)$ from s' only if s did not receive $\mathcal{D}(i)$ before and $h(i, s') < \bar{h}$. If no information is exchanged between s and s' at time k , s stops contacting any of its neighbors. However, s may restart contacting its neighbors again if it updated D_s after time $k + 1$.

When s constructs its D_s , it starts by pruning $\mathcal{D}(s)$. The pruning is performed by deleting a virtual sensor $j \in \mathcal{D}(s)$ (i.e., $\mathcal{D}(s) \leftarrow \mathcal{D}(s) \setminus \{j\}$) if none of the virtual sensors that are connected to j , $\{j' \in V : (j, j') \in E\}$, is not included in any received $\mathcal{D}(i)$, i.e. $j \notin \mathcal{D}(i) : \mathcal{D}(i) \in D_s$. This pruning rule ensures that the virtualized sensors maintain the required topology E and the constructed benefit matrices shall result in a feasible virtualization.

Phase III—Construction of Benefit Matrices

As mentioned earlier, finding a feasible virtualization, $\{A, \mathcal{M}_A\}^*$, that maximizes the total benefit given in Eq. (2) is a hard problem due to the large size of the solution space. Therefore, this phase proposes an efficient way of solving this virtualization problem. Specifically, we propose a method that solves this problem in a distributed manner and without requiring any synchronization among sensors, as described next.

During this phase, each sensor s locally constructs its own set, $A^{(s)}$, of g sensors that s chooses as virtualized sensors to assign to virtual sensors in V . Each sensor s also maintains g row vectors, $B_i^{(s)} \in \mathbf{R}^{1 \times g}$ and $i \in A^{(s)}$, that we define as the benefit vector of sensor i seen by s , where the j -th element, $B_{i,j}^{(s)}$, denotes the benefit of assigning participatory sensor $i \in A^{(s)}$ to the virtual sensor $j \in V$ as seen by s , and is given by

$$B_{i,j}^{(s)} = \begin{cases} \alpha \frac{C(i) - R(j)}{C(i)} + \beta \frac{\bar{h} - h(j, s)}{\bar{h}} & \text{if } j \in \mathcal{D}(i), \\ 0 & \text{otherwise.} \end{cases}$$

Our objective is then to construct, for each $s \in S$, the benefit matrix $B^{(s)} = [B_{i \in A^{(s)}}^{(s)}]$ as fast as possible, and find a feasible virtualization, $\{A, \mathcal{M}_A\}$, that maximizes the total benefit,

$$\sum_{(i,j) \in \mathcal{M}_A} B_{i,j}^{(s)},$$

among all $s \in S$ without knowing the \mathcal{G} structure. Moreover, the path length between a sensor s and any other sensor i that s includes in its benefit matrix must not exceed \bar{h} . Finally, a sensor s shall include only the benefit vectors of the g sensors with the largest possible benefit.

Each sensor s initially sets $A^{(s)} = A^{(s)} \cup \{s\}$ if $\mathcal{D}(s) \neq \emptyset$, sets $h(i, s) = 0$ for all $i \in S$, and sets

$$B_{s,j}^{(s)} = \begin{cases} \alpha \frac{C(s) - R(j)}{C(s)} + \beta, & j \in \mathcal{D}(s), \\ 0, & \text{otherwise.} \end{cases}$$

Also, s maintains a scalar, b_s^{\min} , defined as the minimum total benefit it has received from any other sensor and written as

$$b_s^{\min} = \min_i \sum_{j \in V} B_{i,j}^{(s)},$$

and the corresponding sensor,

$$i_s^{\min} = \operatorname{argmin}_i \sum_{j \in V} B_{i,j}^{(s)}.$$

Initially, $b_s^{\min} = 0$ and remains so until $|A^{(s)}| = g$.

Using the same transition matrix, T , defined in Eq. (4), s contacts its neighbor s' only once at each time k . Then, for all $i \in A^{(s)} : i \neq s'$, s pushes the benefit vector $B_i^{(s)}$ to s' only if $h(i, s) < \bar{h}$ and

$$\sum_{j \in V} \left(B_{i,j}^{(s)} - \frac{\beta}{h} \right) > b_{s'}^{\min}.$$

Also, for all $i \in A^{(s')} : i \neq s$, s pulls the benefit vector $B_i^{(s')}$ from s' only if $h(i, s') < \bar{h}$ and

$$\sum_{j \in V} \left(B_{i,j}^{(s')} - \frac{\beta}{h} \right) > b_s^{\min}.$$

If no information is exchanged between s and s' at time k , s stops contacting its neighbors at time $k+1$. However, s may restart contacting its neighbors again if $B^{(s)}$ is updated after time $k+1$.

When s receives $B_i^{(s')}$, s updates $B_{i,j}^{(s)}$ as

$$B_{i,j}^{(s)} = \begin{cases} B_{i,j}^{(s)} - \frac{\beta}{h} & \text{if } j \in \mathcal{D}(i), \\ 0 & \text{otherwise.} \end{cases}$$

If $i \notin A^{(s)}$, then we have two scenarios. In the first scenario, s still has not received g benefit vectors, so $b_s^{\min} = 0$ and $|A^{(s)}| < g$, then s updates its set of candidate sensors as $A^{(s)} = A^{(s)} \cup \{i\}$. In the other scenario in which $|A^{(s)}| = g$, s replaces the sensor corresponding to the minimum total benefit, i_s^{\min} , with i so that $A^{(s)} = A^{(s)} \setminus \{i_s^{\min}\} \cup \{i\}$. On the other hand, if $i \in A^{(s)}$, then s updates $B_{i,j}^{(s)}$ if $\sum_{j \in V} B_{i,j}^{(s')} > \sum_{j \in V} B_{i,j}^{(s)}$. Finally, s updates b_s^{\min} , i_s^{\min} , and $h(i, s)$ as $h(i, s) = h(i, s') + 1$.

Finding a feasible virtualization that maximizes the benefit $B^{(s)} = [B_{i,j}^{(s)}]_{i \in A^{(s)}}$ instead of the benefit given in Eq. (2) makes the problem easier because every sensor has a different value for the benefit $B_{i,j}$ that depends only on the length of the physical path between i and s instead of the path lengths of all possible combinations of sensor pairs (i, i') that can virtualize a virtual link. Intuitively, this relaxation still leads to an optimal or near optimal virtualization, because for a connected swarm \mathcal{G} , the number of sensors that are directly connected by a single physical link (clique) grows logarithmically in n and hence this number is larger than g almost surely as $g \ll n$. In such a case, it is sufficient to

ensure that the length of the paths between i and s and between i' and s are the shortest possible ones to ensure that the length of the path between i and i' is also the shortest, as in this case, s , i , and i' reside in the same clique with high probability. We evaluate the effectiveness of this relaxation in Section 5 and show that our virtualization algorithm performs well even when the condition $g \ll n$ does not hold.

Phase IV—Solving Local Assignment Problem

After reception of the g benefit vectors, s proceeds to this phase of the algorithm only if it stops communicating and $|A^{(s)}| = g$. Each sensor $s \in S$ with $|A^{(s)}| = g$ solves locally the following assignment problem:

$$\begin{aligned} & \text{maximize} && \sum_{i \in A^{(s)}} \sum_{j \in \mathcal{D}(i)} B_{i,j}^{(s)} m_{ij} \\ & \text{subject to} && \sum_{j \in \mathcal{D}(i)} m_{ij} = 1, \quad i \in A^{(s)}, \\ & && \sum_{\{i: j \in \mathcal{D}(i)\}} m_{ij} = 1, \quad j \in V, \\ & && m_{ij} \in \{0, 1\}, \end{aligned} \tag{5}$$

where m_{ij} are binary optimization variables indicating whether the participatory sensor i is assigned to the virtual sensor j . The problem formulated in (5) is equivalent to the maximum weight matching perfect problem in a bipartite graph, and hence, we propose to use the classical Hungarian method to solve it (the worst case time complexity is $O(g^3)$ [13, 14]).

We can also tolerate an error $\epsilon > 0$ of the resulting total benefit and relax the restriction of finding a perfect matching for large g . This relaxation is reasonable when there are enough sensors involved in solving these local optimization problems, as in this case we can pick the best solution and discard those without a perfect matching. In such a scenario, we can also use a linear time $(1 - \epsilon)$ -approximation algorithm to solve (5) [8]. In this paper, we use the Hungarian method to solve our formulated optimization problems. Details of the algorithm are omitted due to space limitation; readers are referred to [13, 14, 8] for detailed information.

Each sensor solves locally the optimization problem given in (5) and sends its obtained solution to the cloud agent. This is done asynchronously. The cloud agent then selects the solution that leads to the maximum total benefit, and keeps all other solutions for later use in the event that the network dynamics invalidate the selected solution before the virtual sensing task completes.

Complexity and message overhead. The time required to spread Υ across the network is $O(r^{-1} \log n)$ [19]. It takes $O(g)$ worst case time to evaluate $\mathcal{D}(i)$ locally at sensor i . Also, the time required to spread information in the pruning and benefit construction phases is $O(r^{-1} n \log n)$. The pruning of the virtual domain $\mathcal{D}(i)$ requires node i to examine g received virtual domains, each having at most g entries. The worst case local running time of pruning is then $O(g^2)$. Finally, the local running time of the Hungarian method is $O(g^3)$. Hence, the overall complexity of is $O(\max\{r^{-1} n \log n, g^3\})$.

The average number of messages communicated per sensor during the sensor search phase is $\Theta(1)$ and each message is $O(g)$ in size. During pruning of virtual domains, since every sensor exchanges a maximum of n domains each of size that is also $O(g)$, the average number of messages communicated per sensor is $O(n)$. However, because we restrict that messages be communicated up to \bar{h} hops for only a group of sensors that support the requirements of Υ , the average number of messages per sensor is typically small. Figure 1 shows the total time and the average number of messages per sensor required during both the domain pruning and the benefit construction phases. The total time growth is linearithmic in n when Υ is sent to exactly one sensor and when \mathcal{G} is connected. This time can, in practice, be decreased significantly if Υ is initially sent to multiple sensors. Additionally, the average number of messages per sensor is shown to scale linearly with n , and is typically a very small fraction of n .




Figure 1: Time (in number of iterations) and message overhead (in number of communicated messages) resulting from constructing the benefit matrices.

4 Distributed Estimation of θ

After completing the sensor virtualization task, using the proposed RADV algorithm that we described in Section 3, the virtual sensors run an in-network parameter estimation algorithm to compute $\hat{\theta}$ distributedly. In this section, we present our proposed Randomized and Asynchronous Distributed Estimation (RADE) algorithm. We first follow the standard ADMM approach to derive primal, dual and Lagrangian variable update equations, then we describe the proposed RADE algorithm. For clarity of notation, in what follows, we refer to the set of g selected participatory sensors, determined by means of the proposed RADV algorithm, simply as A .

The centralized estimation approach given in (3) is first decomposed into g local estimates of θ (one $\hat{\theta}_i$ for each $i \in A$) while constraining the local estimates with the coupling constraints $\theta_i = \theta_j$ for all $(i, j) \in P$. This results in the following optimization problem:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \sum_{i \in A} \|x_i - H_i \theta_i\|^2 \\ & \text{subject to} && \theta_i - \theta_j = 0 \text{ for all } (i, j) \in P, \end{aligned} \quad (6)$$

where $\{\theta_i, i \in A\}$ are the optimization variables.

By introducing an auxiliary variable, z , we decouple the constraints in (6), so that $\theta_i - z = 0$ for all $i \in A$ [15]. However, this requires that z be shared among all g sensors. Instead, we introduce g auxiliary variables, z_i , and equivalently write the optimization problem as

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \sum_{i \in A} \|x_i - H_i \theta_i\|^2 \\ & \text{subject to} && \theta_j - z_i = 0 \text{ for all } (i, j) \in P. \end{aligned} \quad (7)$$

Let $\lambda = \{\lambda_{i,j} \in \mathbf{R}^{N \times 1} : (i, j) \in P\}$ and $\rho = \{\rho_{i,j} \in \mathbf{R} : (i, j) \in P\}$ denote respectively the set of Lagrangian multipliers and the set of penalty parameters. The augmented Lagrangian is

$$\begin{aligned} \mathcal{L}_\rho(\theta, z, \lambda) = \sum_{i \in A} & \left[\frac{1}{2} \|x_i - H_i \theta_i\|^2 \right. \\ & - \sum_{j \in A: (i,j) \in P} \lambda_{i,j}^\top (\theta_i - z_j) \\ & \left. + \sum_{j \in A: (i,j) \in P} \frac{\rho_{i,j}}{2} \|\theta_i - z_j\|^2 \right]. \end{aligned} \quad (8)$$

By setting the gradient w.r.t θ_i of Eq. (8) to zero and solving for θ_i , we get

$$\begin{aligned} \theta_i &= \left(H_i^\top H_i + \sum_{j \in A: (i,j) \in P} \rho_{i,j} I \right)^{-1} \\ & \cdot \left(H_i^\top x_i + \sum_{j \in A: (i,j) \in P} (\lambda_{i,j} + \rho_{i,j} z_j) \right). \end{aligned}$$

Similarly, we solve for z_i by setting the gradient w.r.t to z_i to zero and rearranging the indices of the Lagrangian multipliers and the penalty parameters. It follows that

$$z_i = \frac{1}{g} \sum_{j \in A: (i,j) \in P} \left(\theta_j - \frac{1}{\rho_{j,i}} \lambda_{j,i} \right).$$

The former analysis leads to the conventional ADMM-based distributed consensus estimation algorithm given by

$$\begin{aligned} \theta_i^{(k+1)} &= \left(H_i^\top H_i + \sum_{j \in A: (i,j) \in P} \rho_{i,j} I \right)^{-1} \\ & \cdot \left(H_i^\top x_i + \sum_{j \in A: (i,j) \in P} \left(\lambda_{i,j}^{(k)} + \rho_{i,j} z_j^{(k)} \right) \right), \\ z_i^{(k+1)} &= \frac{1}{g} \sum_{j \in A: (i,j) \in P} \left(\theta_j^{(k)} - \frac{1}{\rho_{j,i}} \lambda_{j,i}^{(k)} \right), \\ \lambda_{j,i}^{(k+1)} &= \lambda_{j,i}^{(k)} - \rho_{j,i} \left(\theta_j^{(k+1)} - z_i^{(k+1)} \right), \end{aligned} \quad (9)$$

where the superscript k denotes the value of the variable at the k -th iteration. This conventional ADMM algorithm, given in (9), requires synchronization and variable update among the sensors [6, 22]. Moreover, at each iteration k , each sensor i must send $z_i^{(k)}$ and $\theta_i^{(k)}$ to all other sensors it is connected to, so as to evaluate their $k+1$ primal, dual, and Lagrangian multipliers. When M is small, this algorithm incurs communication overhead that can be shown

to be worse than the communication overhead incurred by centralized estimation methods. However, when M is large, the conventional ADMM algorithm incurs lesser communication overhead than what centralized estimation methods incur, but it still remains practically unattractive due to other weaknesses, detailed later in Section 5.

Given the absolute and relative tolerances, ϵ_{abs} and ϵ_{rel} , specified by the SLAs between the cloud client and the cloud agent, we define the primal and dual tolerances, controlling the convergence of the algorithm at iteration k , as

$$\epsilon_i^{\text{pri}}(k) = \sqrt{g}\epsilon_{\text{abs}} + \epsilon_{\text{rel}} \max(\|\theta_i^{(k)}\|, \|z_i^{(k)}\|),$$

and

$$\epsilon_i^{\text{dual}}(k) = \sqrt{g}\epsilon_{\text{abs}} + \epsilon_{\text{rel}} \sum_{j \in A} \|\rho_{j,i} \lambda_{j,i}\|.$$

The tolerances, ϵ_i^{pri} and ϵ_i^{dual} , define the stopping criteria of sensor i ; i.e., sensor i stops updating θ_i and z_i when

$$\|\theta_i^{(k+1)} - z_i^{(k+1)}\| < \epsilon_i^{\text{pri}}(k), \quad (10)$$

and

$$\|z_i^{(k+1)} - z_i^{(k)}\| < \epsilon_i^{\text{dual}}(k). \quad (11)$$

The stopping criteria of RADE are different from those of the conventional ADMM. Unlike the conventional ADMM where all sensors shall stop computations all at the same time using a common stopping criteria and common primal and dual tolerances, the stopping criteria (10) and (11) of RADE allow a sensor i to stop its computations asynchronously and independently from other sensors. However, these criteria are not enough to ensure asynchronous implementation, as synchronization is still required for dual and primal variable updates at iteration $k+1$ due to their dependencies on k .

To ensure full asynchronous implementation, we use the doubly stochastic transition matrix, $T \in \mathbf{R}^{g \times g}$, where $T_{i,j}$ is the probability that a sensor i contacts another sensor j at any iteration, for deciding the communications among sensors. We can have

$$T_{i,j} = \begin{cases} \frac{1}{d'_i + 1} & \text{if } i = j \text{ or } (i,j) \in P, \\ 0 & \text{otherwise,} \end{cases}$$

where $d'_i = |\{j \in A : (i,j) \in P\}|$ is the degree of the virtual sensor, in Υ , that i virtualizes. At iteration $k+1$, sensor i may need to contact only one sensor j , unless both of i 's stopping criteria, (10) and (11), are already satisfied. Whereas sensor j can be contacted by more than one sensor if j is not contacting any other sensor, even when both of j 's stopping criteria are satisfied.

Upon contacting j , sensor i pushes $\theta_i^{(k)}$ to j only if i 's primal stopping condition is not satisfied and pushes $z_i^{(k)}$ to j only if i 's dual stopping condition is not satisfied. Also, i pulls $\theta_j^{(k)}$ from j only if j 's primal stopping condition is not satisfied and pulls $z_j^{(k)}$ only if j 's dual stopping condition is not satisfied. Finally, both i and j update their $k+1$




Figure 2: MSE of RADE compared to those achieved under ADMM and LS at different noise power and for different virtual sensor network topologies.

variables using the most recent values they received from other sensors.

Mean square error and convergence. The asynchronous and randomization design features of RADE do not impact the Mean Square Error (MSE) achieved by RADE when compared to ADMM. This is explained as follows. In both ADMM and RADE, the number of necessary dual and primal variables updates required until convergence remains unchanged, so that convergence to the same estimate is guaranteed in both algorithms. Figure 2 shows the MSE achievable under both RADE and ADMM when compared to LS under each of the three studied sensor network topologies: complete, star, and cycle. These results show the optimality of RADE that we intuitively discussed. All approaches have the same accuracy. But of course each of them does so at a different performance cost, as will be discussed later.

On the other hand, RADE exhibits a linear convergence rate ($O(1/k)$), similar to what the conventional ADMM does. Figure 3 shows the number of time steps required for both RADE and ADMM to converge under different relative tolerance parameters, ϵ_{rel} . RADE convergence tends to be more restricted by the randomization nature of the algorithm for smaller values of ϵ_{rel} , which can be seen by the increasing number of steps as g increases if $\epsilon_{\text{rel}} = 10^{-2}$. ADMM generally requires a lesser number of steps to converge by relaxing the consensus constraint (through reducing ϵ_{rel}). However, as will be seen in the numerical results section later, this increase in the number of convergence steps is acceptable when considering the amount of communication overhead that the algorithm saves.

5 Numerical Results

In this section, we evaluate the performance of the proposed RADV and RADE algorithms through simulations. In our simulations, the swarm of sensors, \mathcal{G} , and the virtual sensing task requests, Υ , are generated using the parameters summarized in Table 1. The virtual sensor network topology can either be complete, cyclic, or star, with a randomly chosen central location, c . We consider receiving and servicing only one virtual sensing task request at a time. The




Figure 3: Number of time steps (k) needed until convergence of RADE when compared to ADMM for complete topology under different relative tolerance values ϵ_{rel} .




Figure 4: Rejection rate encountered at different n .

absolute and relative tolerances, ϵ_{abs} and ϵ_{rel} , are set to 10^{-4} unless specified otherwise.

Table 1: Simulation Parameters

Parameter	r	$C(i)$	$R(j)$	δ	\bar{h}
Value	0.1	$\sim U(50, 100)$	$\sim U(25, 50)$	0.2	20

Figure 4 shows the rejection rate encountered with different Υ topologies and n values. As we only consider one single request at a time, the results shown in this figure reflect mainly the impact of the virtual sensor network topology, the number of sensors n , and the simulations parameters given in Table 1 on the rejection rate. The denser the swarm of sensors is, the lower the rejection rate is, implying that the cloud is capable of granting higher number of requests. As a marginal note, a star topology is slightly easier to virtualize than a complete or a cyclic topology.

One way of assessing the effectiveness of the virtualization algorithm is by measuring the difference between the total virtualization benefit given in (2) and the cost associated with the sensor virtualization introduced in Section 2.2.2. For a given number of virtual sensors, the cost is mainly determined by the choice of the topology (star topology has the lowest cost and complete topology has the highest one). For a given topology, the total benefit is maximized when each virtual sensor is assigned to the participatory sensor with the maximum capacity and each virtual link is mapped to exactly one physical link. We refer to this maximized benefit as the upper bound.




Figure 5: Virtualization cost of RADV when compared to the upper bound under different topologies.




Figure 6: Number of time steps until convergence of RADE when compared to ADMM under different topologies.

In Figure 5, we evaluate the virtualization effectiveness achieved by RADV under different virtual topologies. As the swarm gets denser, RADV achieves a Total Benefit – Cost that is very close to the upper bound. Since the lowest possible virtualization cost is with star or cyclic topologies, it is desired by the cloud to always arrange each virtual sensing task in a star or a cyclic topology. On the other hand, convergence and communication overhead of the distributed estimation is also impacted by the cloud agent's choice of the virtual topology. This creates a design trade-off, as we will see in the next two paragraphs.

Figure 6 shows the impact of the virtual topology choice on the convergence performance of RADE when compared to ADMM. If g is small (three to eight), the impact of the virtual topology on convergence of RADE and ADMM is minimal. This is because the degree of parallelism (number of sensors active at the same time) is more restricted by the small number of virtual sensors g . In such a scenario, it is convenient for the cloud agent to always arrange the virtual sensors in a star topology. However, as g increases, the impact of the virtual topology becomes significant as the degree of parallelism is higher in a complete topology, enabling RADE to converge much faster as g gets larger. This convergence becomes slower with star and cyclic topologies. This is because in star and cyclic topologies, only few sensors are active at a time, making RADE and ADMM converge in a number of steps comparable to that of the ADMM's sequential implementation. In this later scenario, the cloud agent shall arrange the virtual sensors as a complete topol-




Figure 7: Communication overhead when comparing RADE to ADMM and LS under different g values.

ogy unless the SLA permits slower convergence.

Moreover, RADE converges in a higher number of steps when compared to the conventional ADMM. This is because in ADMM, all sensors are active at each time, and a sensor exchanges its updated variables with all of its neighbors, whereas in RADE, only disjoint sensor pairs are active at a time and variables are updated only between pairs of sensors. Nevertheless, we argue that this loss in speed of convergence for RADE is marginal when compared to the significant savings in communication overhead.

Figure 7 shows the total number of $O(N)$ sized messages exchanged during estimation when comparing RADE, ADMM, and LS for $M = 100$. The number of messages exchanged by RADE is at least an order of magnitude less than the number of messages generated under ADMM. Also the communication overhead of RADE is less than the centralized LS especially as M becomes large. This savings in communication overhead is attributed to the asynchronous design of RADE in which messages among sensors are only exchanged if new values of a primal or dual variables are changed away from their specified tolerances.

6 Conclusion and Discussion

We propose cloud-based remote sensing algorithms for enabling distributed estimation of unknown parameters via sensor network virtualization. The algorithm has the following phases: sensor search, domain pruning, benefit matrix construction, virtual-participatory sensor assignment solver, and distributed estimation. Using simulation, we show that the proposed algorithms reduce communication overhead significantly without compromising the estimation error when compared to the traditional ADMM algorithm. We also show that the convergence time of our proposed algorithms maintain linear convergence behavior, as in the case of conventional ADMM.

References

[1] S. Abdelwahab, B. Hamdaoui, M. Guizani, and A. Rayes. Enabling smart cloud services through re-

mote sensing: An internet of everything enabler. *Internet of Things Journal, IEEE*, 1(3):276–288, June 2014.

- [2] M. T. Beck, A. Fischer, H. de Meer, J. F. Botero, and X. Hesselbach. A distributed, parallel, and generic virtual network embedding framework. In *Communications (ICC), 2013 IEEE International Conference on*, pages 3471–3475. IEEE, 2013.
- [3] A. Belbekkouch, M. Hasan, A. Karmouch, et al. Resource discovery and allocation in network virtualization. *Communications Surveys & Tutorials, IEEE*, 14(4):1114–1128, 2012.
- [4] K. Birman. The promise, and limitations, of gossip protocols. *ACM SIGOPS Operating Systems Review*, 41(5):8–13, 2007.
- [5] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *Information Theory, IEEE Transactions on*, 52(6):2508–2530, 2006.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. 2011.
- [7] X. Cheng, S. Su, Z. Zhang, H. Wang, F. Yang, Y. Luo, and J. Wang. Virtual network embedding through topology-aware node ranking. *ACM SIGCOMM Computer Communication Review*, 41(2):38–47, 2011.
- [8] R. Duan and S. Pettie. Linear-time approximation for maximum weight matching. *Journal of the ACM (JACM)*, 61(1):1, 2014.
- [9] S. Ferretti. Gossiping for resource discovering: An analysis based on complex network theory. *Future Generation Computer Systems*, 29(6):1631–1644, 2013.
- [10] A. Fischer, J. F. Botero, M. Till Beck, H. De Meer, and X. Hesselbach. Virtual network embedding: A survey. *Communications Surveys & Tutorials, IEEE*, 15(4):1888–1906, 2013.
- [11] F. Garin and L. Schenato. A survey on distributed estimation and control applications using linear consensus algorithms. In *Networked Control Systems*, pages 75–107. Springer, 2010.
- [12] I. Houidi, W. Louati, and D. Zeghlache. A distributed virtual network mapping algorithm. In *Communications, 2008. ICC’08. IEEE International Conference on*, pages 5634–5640. IEEE, 2008.
- [13] H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [14] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial & Applied Mathematics*, 5(1):32–38, 1957.

- [15] D. P. Palomar and M. Chiang. A tutorial on decomposition methods for network utility maximization. *Selected Areas in Communications, IEEE Journal on*, 24(8):1439–1451, 2006.
- [16] H. Paul, J. Fliege, and A. Dekorsy. In-network-processing: Distributed consensus-based linear estimation. *Communications Letters, IEEE*, 17(1):59–62, 2013.
- [17] C. Perera, P. P. Jayaraman, A. Zaslavsky, D. Georgakopoulos, and P. Christen. Sensor discovery and configuration framework for the internet of things paradigm. In *Internet of Things (WF-IoT), 2014 IEEE World Forum on*, pages 94–99. IEEE, 2014.
- [18] C. Perera, A. Zaslavsky, P. Christen, M. Compton, and D. Georgakopoulos. Context-aware sensor search, selection and ranking model for internet of things middleware. In *Mobile Data Management (MDM), 2013 IEEE 14th International Conference on*, volume 1, pages 314–322. IEEE, 2013.
- [19] D. Shah. *Gossip algorithms*. Now Publishers Inc, 2009.
- [20] P. Tarrío, A. M. Bernardos, and J. R. Casar. Weighted least squares techniques for improved received signal strength based localization. *Sensors*, 11(9):8569–8592, 2011.
- [21] E. Wei and A. Ozdaglar. On the $o(1/k)$ convergence of asynchronous distributed alternating direction method of multipliers, 2013.
- [22] G. Xu, H. Paul, D. Wuebben, and A. Dekorsy. Fast distributed consensus-based estimation (fast-dice) for cooperative networks. In *Smart Antennas (WSA), 2014 18th International ITG Workshop on*. VDE, 2014.